

Fixing (?) the Shortcomings of Map-based Approaches



MAP OF THE INTERNET
THE IPv4 SPACE, 2006
<http://xkcd.com/195/>

Provisioned Bandwidth Information

- Applications may want to use it
 - Ono and P4P experiments show higher improvement in high-capacity areas
- ISPs have it and may want applications to use it
 - “Try local power-users before going to Korea!”
- BTW, it's in the charter...

Problem (Part 1)

- IP of residential customers primarily assigned on a topological basis
 - In the same area addresses often taken from the same pool regardless of the subscription class

i.e.

- Prefix matching **inefficient** for identifying access line characteristics
 - Cannot tell between 50/10 and 2/.384 lines

Example (Priority Maps)

Topology only

```
83.128.0.0/10 → 20 /* Local AS */
83.128.0.0/12 → 30 /* Local PoP */
85.128.0.0/9 → 4
91.0.0.0/10 → 10 /* Peering agr */
99.98.0.0/16 → 5
202.0.0.0/8 → 0 /* Australia!!! */
204.2.0.0/16 → 7
...
```

Topology and bandwidth

```
83.123.21.45/32 → 30 /* 50/10 */
83.123.21.46/32 → 20 /* 20/2*/
83.123.21.47/32 → 1 /* Dialup */
83.123.21.48/31 → 30 /* 2 in a row! */
83.123.21.50/32 → 10 /* 10/1 */
83.123.21.50/32 → 30 /* 50/10 */
83.123.21.51/32 → 20 /* 20/2 */
...
```

Problem (Part 2)

- IP addresses of residential lines often assigned dynamically

i.e.

- Address-by-address maps **stale** as soon as the first home gateway reboots

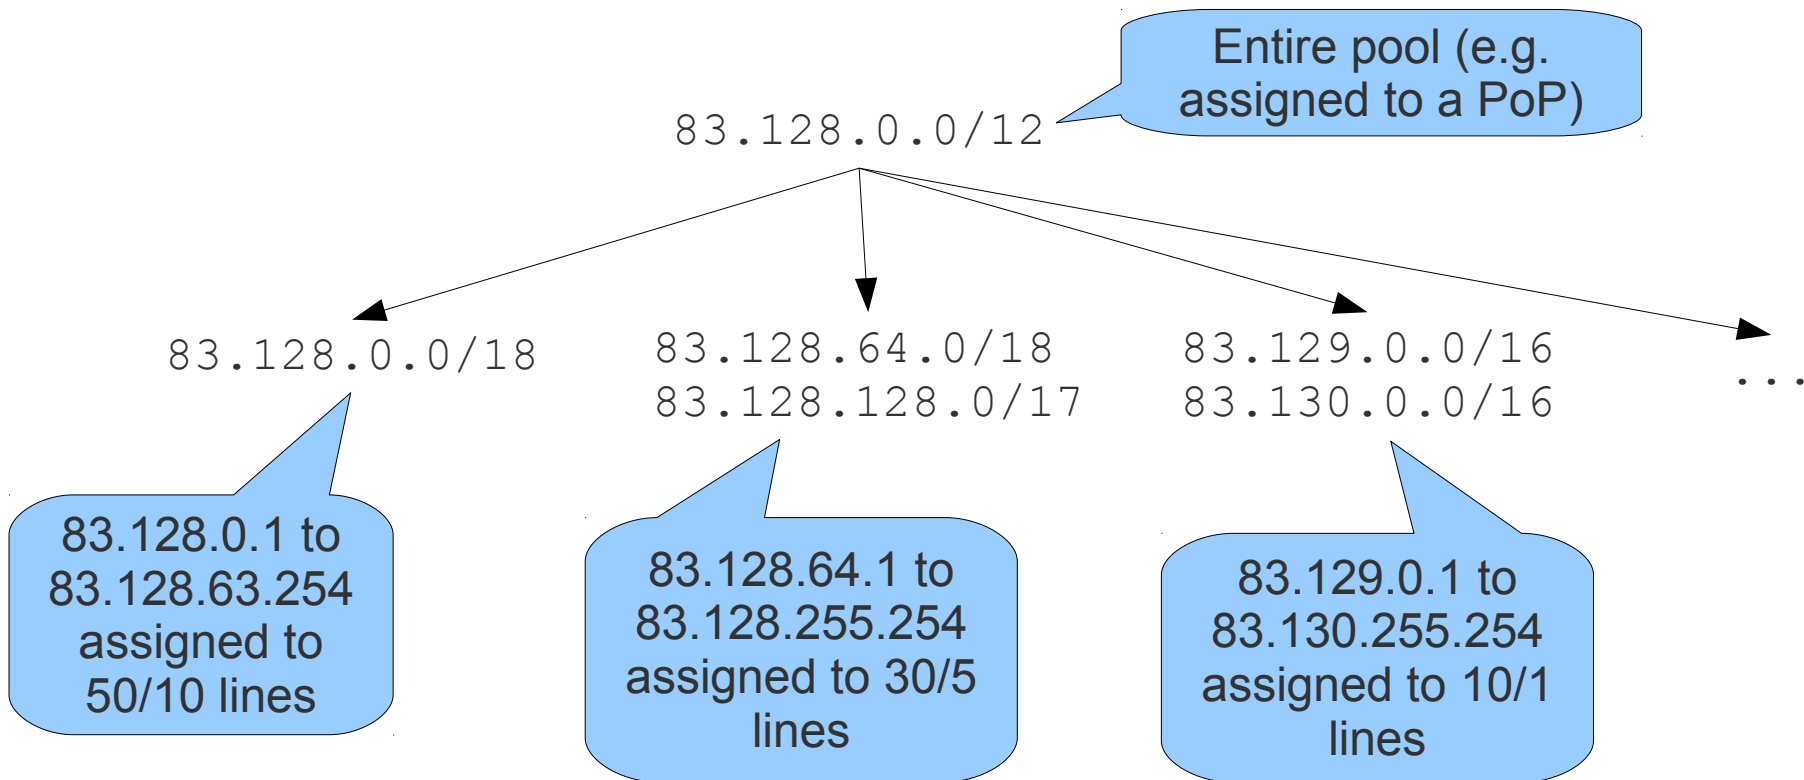
Question: Is it *desirable* to be able to provide guidance based on provisioned bandwidth? (Or on anything other than topology?)



Suggested answer: Yes, if reasonably feasible

- Huge, short-lived maps may turn out generating more traffic than the application itself
- Applications may not want to give privacy in exchange
BitTorrent won't send lists of IP addresses to ISPs, no matter how accurately they can do the ranking

Approach #1

- **Adapting IP provisioning policies to reflect provisioned bandwidth**
 - Multi-level network partitioning





Approach #1: Pros and Cons

Pros


- Still a map
 - Simple, simplest

Cons

- May require radical changes in ISPs' provisioning policies
- Additional partitioning unlikely to reflect different parameters (bandwidth, latency...)
- May result in big files

Approach #2



- Fine-grain guidance in an additional step
 - First: a high-level map
 - “Prefer Korea, avoid Australia...”
 - Then: [take a deep breath] an accurate query/response service for areas where more detailed information is available
- Steps *almost* completely disjointed
 - Maps may point to second-step servers
 - “Prefer Korea, avoid Australia, detailed information about Japan available at xyzp://oracle.isp.jp...”




Peer List




Preference Map




Coarse-grain Refined List





Oracle Location(s)



Coarse-grain Refined List



Oracle(s)



Fine-grain Refined List

Approach #2: Pros and Cons

Pros

- No impact on ISP infrastructure/policies
- Incremental solution
 - Not a one-size-fits-all
 - Can stop at step 1

Cons

- Added complexity
 - Discovery of maps and oracles
 - One protocol or two protocols?
- Privacy
 - Not all users will be happy with step 2

Approach #3

- Click to add your favorite solution

We don't need answers now, but...

At some point we'll have basically two options:

- Ignore everything that does not fit on a map
 - Provisioned bandwidth is not available bandwidth
 - IP assignment policies are just policies
 - Just topology is still better-than-random
- Design alternatives/complements to prefix-matching
 - Efficiency (no, /30 is not prefix-matching!)